arXiv:1407.5509v1 [math.ST] 21 Jul 2014

Inconsistent treatment estimates from mis-specified logistic regression analyses of randomized trials

J.N.S. Matthews *

School of Mathematics & Statistics, Newcastle University, Newcastle upon Tyne, NE1 7RU, UK

Nuri H. Badi †

Statistics Department, University of Benghazi, Benghazi, Libya

Abstract

When the difference between treatments in a clinical trial is estimated by a difference in means, then it is well known that randomization ensures unbiassed estimation, even if no account is taken of important baseline covariates. However, when the treatment effect is assessed by other summaries, e.g. by an odds ratio if the outcome is binary, then bias can arise if some covariates are omitted, regardless of the use of randomization for treatment allocation or the size of the trial. We present accurate closed-form approximations for this asymptotic bias when important Normally distributed covariates are omitted from a logistic regression. We compare this approximation with ones in the literature and derive more convenient forms for some of these existing results. The expressions give insight into the form of the bias, which simulations show is usable for distributions other than the Normal. The key result applies even when there are additional binary covariates in the model.

Key words *Asymptotic bias; baseline values; logistic regression; probit regression; randomized clinical trial.*

*john.matthews@ncl.ac.uk: corresponding author

†nhb_2020@yahoo.com

1 Introduction

Randomized trials are often analysed using a linear or generalized linear model, so that the treatment effect can be adjusted for important baseline covariates. However, if some baseline variables cannot be measured, or if their importance is not appreciated, then they will be omitted from the model. Randomization ensures that the estimate of the treatment effect is unbiassed when relevant covariates are omitted from a linear model. This is a consequence of the unit-treatment additivity in such models [1, chapter 5] and does not necessarily carry over to generalized linear models. Several non-linear models for which unbiassed estimators are obtained, notwithstanding the omission of covariates, are identified in [2], who also show that the important case of binary outcomes analysed using a logistic model is asymptotically biassed when covariates are omitted.

Numerous authors have addressed the problem of the effect of the omission of covariates in logistic regression. In biostatistical contributions an epidemiological perspective is, perhaps more common [3, 4, 5, 6, 7, 8], some authors do focus on randomized trials [2, 9, 10, 11]. Gail and colleagues [2] derive approximations for the asymptotic bias in the treatment estimator when all covariates other than the treatment indicator are omitted. The case of two general, scalar, covariates, one of which is fitted and the other omitted is considered in [10, 8]. The main exposition in [10] assumes that the covariates are independent but, as the authors explain, this restriction can be relaxed. In all these articles Taylor series approximations are used to provide some indication of the size and direction of the bias, so the expressions derived are necessarily restricted to small parameter values, although whether it is the parameter of the fitted or omitted covariate that needs to be small varies between these contributions.

In this article we make use of the properties of the extended skew-normal distribution [12] and an approximation of the logistic function by the probit to obtain expressions for the least false values [13, p.25] of the fitted covariates when other covariates are omitted. No use of Taylor series approximations is required, so the expressions give excellent numerical results for a wide range of parameter values and provide useful insight into the form the bias takes in a randomized trial. Our main result applies to a logistic regression with a single binary covariate, which we usually take to indicate the treatment allocation, and an arbitrary number of continuous covariates. The latter are assumed to follow a multivariate Normal distribution but simulation results show that the results hold for a wider class of covariates. Explicit forms for the asymptotic bias given in [9, 10] are derived for our case and compared with that found using the skew-Normal distribution. Extensions to allow

additional binary covariates is possible, although these extensions require further assumptions.

In the next section we present the expression for the least false values and in Section 3 related work is explored. Some simulation results are given in Section 4, extensions to allow additional binary covariates are discussed in Section 5 and some conclusions are drawn in the final section.

2 Least false values

Suppose that the random variable $Y \in \{0, 1\}$ is related to a binary covariate $T \in \{-1, 1\}$ and further covariates X_1 and X_2, that have p and q dimensions respectively, by

$$\Pr(Y = 1 \mid T, X_1, X_2) = \text{expit}(\mu + \alpha T + \beta_1^T X_1 + \beta_2^T X_2) \tag{1}$$

where $\text{expit}(u) = \exp(u)/[1 + \exp(u)]$. If the fitted model omits X_2, i.e. if

$$\Pr(Y = 1 \mid T, X_1) = \text{expit}(\mu + \alpha T + \beta_1^T X_1) \tag{2}$$

is assumed to apply then, as the sample size increases, the maximum likelihood estimates $(\hat{\mu}, \hat{\alpha}, \hat{\beta}_1)$ will tend to the least false values $(\mu^*, \alpha^*, \beta_1^*)$. From the score equations it can be shown [13, p25] that

$$\begin{aligned}
\mathsf{E}[\text{expit}(\mu^* + \alpha^* T + \beta_1^{*T} X_1)] &= \mathsf{E}[\text{expit}(\mu + \alpha T + \beta_1^T X_1 + \beta_2^T X_2)] & (3)\\
\mathsf{E}[T\text{expit}(\mu^* + \alpha^* T + \beta_1^{*T} X_1)] &= \mathsf{E}[T\text{expit}(\mu + \alpha T + \beta_1^T X_1 + \beta_2^T X_2)] & (4)\\
\mathsf{E}[X_{1j}\text{expit}(\mu^* + \alpha^* T + \beta_1^{*T} X_1)] &= \mathsf{E}[X_{1j}\text{expit}(\mu + \alpha T + \beta_1^T X_1 + \beta_2^T X_2)] & (5)
\end{aligned}$$

where X_{1j} is the j^{th} element of X_1, $j = 1, \ldots, p$ and expectations are taken with respect to the joint distribution of (T, X_1, X_2).

The density of an extended multivariate skew-Normal (ESN) random variable $U \in \mathbb{R}^p$ [12] is

$$f(u) = \frac{\phi_p(u; \omega, \Omega)\Phi(\zeta^T(u - \omega) + \psi)}{\Phi(\psi/\sqrt{1 + \zeta^T \Omega \zeta})}, \tag{6}$$

where ζ is a p-dimensional parameter, ψ is a scalar, $\phi_p(\cdot; \omega, \Omega)$ is the p-dimensional multivariate Normal density with mean ω and dispersion Ω and $\Phi(\cdot)$ is the standard Normal distribution function. The mean of the ESN distribution is

$$\mathsf{E}(U) = \omega + \frac{\Omega\zeta}{\sqrt{1 + \zeta^T \Omega \zeta}} \frac{\phi(\bar{\psi})}{\Phi(\bar{\psi})},$$

where $\bar{\psi} = \psi(1 + \zeta^T \Omega \zeta)^{-\frac{1}{2}}$ and $\phi(\cdot) = \phi_1(\cdot; 0, 1)$.

We consider the case when, conditional on $T = t$, $X = (X_1^T, X_2^T)^T$ follows a multivariate Normal distribution with mean ν_t and dispersion Ω, $t = -1, 1$. In principle we could allow the dispersion to change with T but analytic progress does not seem possible in this case. We also use $\nu_{t,1}$, $\nu_{t,2}$, Ω_{11}, Ω_{22}, Ω_{12} and Ω_{21} to denote the partition of ν_t and Ω induced by the partition of X. If we use the approximation $\text{expit}(u) \approx \Phi(cu)$ with $c = 16\sqrt{3}/(15\pi)$ [14] in (3), (4) and (5), then properties of the ESN distribution provide the approximations

$$\beta_1^* \approx \frac{\beta_1 + \Omega_{11}^{-1}\Omega_{12}\beta_2}{\sqrt{1 + c^2\beta_2^T\tilde{\Omega}\beta_2}} \tag{7}$$

$$\mu^* \approx \frac{1}{\sqrt{1 + c^2\beta_2^T\tilde{\Omega}\beta_2}}[\mu + \tfrac{1}{2}\beta_2^T\{(\nu_{1,2} + \nu_{-1,2}) - \Omega_{21}\Omega_{11}^{-1}(\nu_{1,1} + \nu_{-1,1})\}] \tag{8}$$

$$\alpha^* \approx \frac{1}{\sqrt{1 + c^2\beta_2^T\tilde{\Omega}\beta_2}}[\alpha + \tfrac{1}{2}\beta_2^T\{(\nu_{1,2} - \nu_{-1,2}) - \Omega_{21}\Omega_{11}^{-1}(\nu_{1,1} - \nu_{-1,1})\}] \tag{9}$$

where $\tilde{\Omega} = \Omega_{22} - \Omega_{21}\Omega_{11}^{-1}\Omega_{12}$, is the dispersion of X_2 conditional on X_1. Outline details of the derivation can be found in the Appendix. Note that if $\tilde{q} = \sqrt{1 + c^2\beta_2^T\tilde{\Omega}\beta_2}$ then the presence of the factor $\tilde{q}^{-1}$ means that even if X_1 and X_2 are uncorrelated, then, unlike the linear model, $\beta_1^* \neq \beta_1$ unless, trivially, $\beta_2 = 0$, or $\tilde{\Omega} = 0$, i.e. the variation in the omitted variables is wholly explained by the fitted variables. To repeat, the only approximation required for the results in (7), (8) and (9) is that of a logistic by a probit, which is well known to be highly accurate.

When T is the treatment indicator and X are baseline covariates from a randomized trial, then the assumption made above, namely $\mathsf{var}(X \mid T = 1) = \mathsf{var}(X \mid T = -1)$ is automatically satisfied and, additionally $\nu_1 = \nu_{-1}$, so (9) implies that the least false value of the treatment effect α is

$$\alpha^* \approx \frac{\alpha}{\sqrt{1 + c^2\beta_2^T\tilde{\Omega}\beta_2}}. \tag{10}$$

Apart from the cases already mentioned which give $\tilde{q} = 1$, (10) shows that the omission of relevant covariates means that the treatment estimator will be biassed towards no effect.

3 Relation with other work

3.1 No fitted covariates other than the treatment indicator

Gail and colleagues[2] considered the bias of treatment estimates for the case when there are no fitted covariates, i.e. the fitted equation is simply

$$\Pr(Y = 1 \mid T) = \text{expit}(\mu + \alpha T),$$

as opposed to (2), and where the omitted covariates are not restricted to being Normally distributed. Finding μ^* and α^* amounts to solving equations (3) and (4) with X_1 omitted. In [2] Taylor series expansions for small $\beta_2^T X_2$ were used to obtain the approximate solution

$$\alpha^* - \alpha \approx -\tfrac{1}{2}\beta_2^T \Omega_{22} \beta_2 (\text{expit}(\mu + \alpha) - \text{expit}(\mu - \alpha)). \tag{11}$$

For small α this is approximately $-\alpha\beta_2^T\Omega_{22}\beta_2\text{expit}(\mu)(1-\text{expit}(\mu))$, whereas (10) implies that for small $\beta_2^T X_2$ the bias is approximately $-\frac{c^2}{2}\alpha\beta_2^T\Omega_{22}\beta_2$. As $\frac{1}{2}c^2 \approx 0.173$, this is similar to $\text{expit}(\mu)(1 - \text{expit}(\mu))$, which varies from 0.1 to 0.25 as μ varies over (-2,2).

In [9, 10] a different approach was applied to the case when the true model has two scalar covariates, only one of which is included in the fitted model. As in [2] no assumption of Normality was made. These authors also used a Taylor series expansion but now applied to the fitted, rather than the omitted covariate. Using the notation in the present paper, and taking T to be the fitted covariate, the approach in [10] noted that

$$\alpha^* = \tfrac{1}{2}[\text{logit}(\pi_1^*) - \text{logit}(\pi_{-1}^*)] = H(\alpha) \tag{12}$$

where $\pi_k^* = \mathsf{E}(\text{expit}(\mu + k\alpha + \xi))$, where the expectation is taken with respect to the distribution of $\xi = \beta_2^T X_2$. Strictly it is the distribution of ξ conditional on $T = k$ but as T is a randomization indicator, this coincides with the unconditional distribution of ξ. Expanding $H(.)$ about $\alpha = 0$ [10] gives, in the case of logistic regression,

$$\alpha^* \approx \alpha H'(0) = \alpha \, \frac{\pi_0^* - \mathsf{E}[\text{expit}(\mu + \xi)^2]}{\pi_0^* - \pi_0^{*2}}, \tag{13}$$

which, as with (10), is seen to be closer to 0 than α. Exact analytic evaluation of $H'(0)$ is not possible but further use of the approximation $\text{expit}(u) \approx \Phi(cu)$ and results due to DB Owen reproduced in [15, p.236], allow (13) to be written as

$$\alpha^* \approx \alpha \frac{2T(h,a)}{\Phi(h)\Phi(-h)} = \alpha \frac{T(h,a)}{T(h,1)} \tag{14}$$

where $h = c(\mu + \beta_2^T \nu_{,2})/\sqrt{1 + c^2 \beta_2^T \Omega_{22} \beta_2}$, $a = 1/\sqrt{1 + 2c^2 \beta_2^T \Omega_{22} \beta_2}$ and $\nu_{,2}$ is the mean of X_2. In (14) $T(h, a)$ is Owen's T function [16], defined as

$$T(h, a) = \frac{1}{2\pi} \int_0^a \frac{\exp[-\frac{1}{2}h^2(1 + x^2)]}{1 + x^2} dx,$$

which has an important role in the computation of bivariate Normal probabilities. It can be evaluated conveniently by the function `T.Owen` in the R package `sn` [17]. For fixed h, $T(h, a)$ is an increasing function of its second argument and as, in the present application, $0 < a < 1$, it follows that the expression for α^* in (14) is always closer to 0 than α. For fixed a, $T(h, a)/T(h, 1)$ is an even function of h and increases as the magnitude of h increases, so the largest attentuation of α occurs at $h = 0$. As the magnitude of h increases, $T(h, a)/T(h, 1)$ approaches one and α^* approaches α.

While (10) gives a bias in α that does not change with the mean of the covariates, this is not the case with (14). This is most accessibly shown by plotting, for a series of values of $\tilde{q}^{-1}$, $T(h, a)/T(h, 1)$ against $P = \text{expit}(\mu + \beta_2^T \nu_{,2})$, which is a typical reponse probability. For most randomized trials P will be between 0.1 and 0.9. The figure shows that the bias correction using (10) is slightly conservative relative to (14) for most values of P. For more extreme P, the bias from $T(h, a)/T(h, 1)$ reduces, as predicted from the behaviour of this expression for larger $|h|$.

3.2 Covariates fitted in addition to the treatment indicator

The approach taken in [10], unlike that in [2], can be adapted to the case when the fitted model includes covariates X_1 in addition to the treatment indicator. For any given X_1 (12) still applies, but with the expectation in $\mathsf{E}[\text{expit}(\mu + k\alpha + \beta_1^T X_1 + \xi)]$ now taken with respect to the distribution of X_2 given X_1. Consequently the bias factor $T(h, a)/T(h, 1)$ still applies but with $a = 1/\sqrt{1 + 2c^2 \beta_2^T \tilde{\Omega} \beta_2}$ and

$$h = \frac{c[\mu + \beta_1^T X_1 + \beta_2^T (\nu_{,2} + \Omega_{21} \Omega_{11}^{-1} (X_1 - \nu_{,1}))]}{\sqrt{1 + c^2 \beta_2^T \tilde{\Omega} \beta_2}}. \tag{15}$$

This is of limited use because of the dependence on X_1, but replacing X_1 by its mean $\nu_{,1}$, so $h = c(\mu + \beta^T \nu)/\sqrt{1 + c^2 \beta_2^T \tilde{\Omega} \beta_2}$, provides a workable alternative that can be compared with (10) when multiple covariates are fitted.

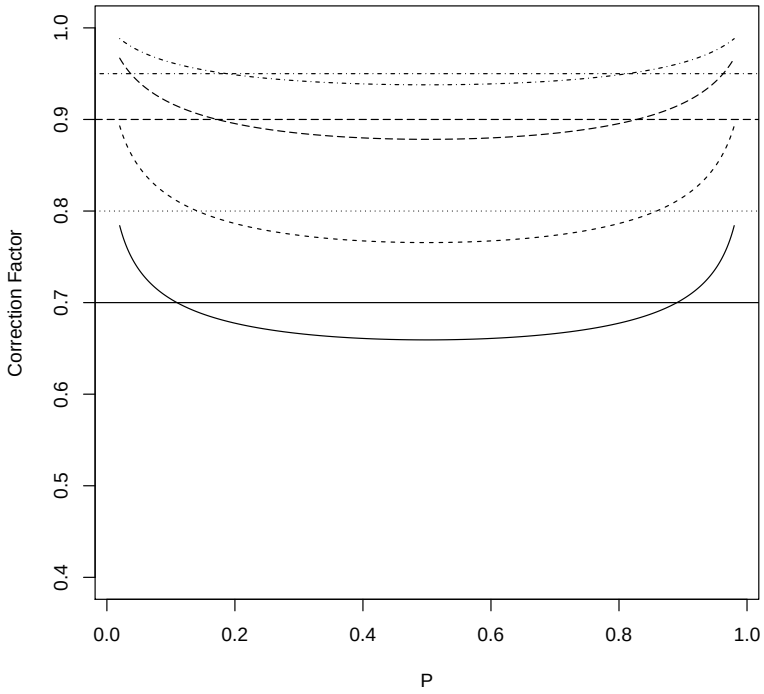

Figure 1: Correction factor $T(h,a)/T(h,1)$ plotted against $P = \text{expit}(\mu + \beta_2^T \nu_{,2})$, for four alternative values of the correction factor $\tilde{q}^{-1}$, namely 0.7 (solid line); 0.8 (dashed line); 0.9 (long-dashed line); 0.95 (dot-dash line). The horizontal lines are at the values of $\tilde{q}^{-1}$

3.3 Probit regression

It is widely acknowledged that in practice logistic and probit regressions can seldom be distinguished in terms of their fit to the data. As the present analyses have exploited the similarity of $\text{expit}(u)$ and $\Phi(u)$ it is natural to consider the use of probit regression as an alternative to logistic regression, i.e. to replace (1) and (2) with $\Pr(Y = 1 \mid T, X_1, X_2) = \Phi(\mu + \alpha T + \beta_1^T X_1 + \beta_2^T X_2)$ etc. The least false values for the maximum likelihood estimators from a probit regression are essentially those in (7), (8) and (9), but with denominator $\sqrt{1 + \beta_2^T \tilde{\Omega} \beta_2}$ in place of $\sqrt{1 + c^2 \beta_2^T \tilde{\Omega} \beta_2}$, although the justification of this result is slightly different - see the Appendix for details. Consequently

$$\alpha^* = \frac{\alpha}{\sqrt{1 + \beta_2^T \tilde{\Omega} \beta_2}} \tag{16}$$

is an exact expression for the asymptotic bias in the treatment effect from a probit regression with a treatment indicator and Normal covariates.

Probit regression was also considered in [2] and [10]. The probit version of (11) is $\alpha^* \approx \alpha(1 - \frac{1}{2}\beta_2^T \Omega_{22} \beta_2)$. The bias term given in [10] for $H'(0)$ for the probit case is $\mathsf{E}(\phi[\Phi^{-1}(\pi_0)])/\phi(\Phi^{-1}[\mathsf{E}(\pi_0)])$. If the true model includes both fitted X_1 and omitted X_2 then the probit analogue of (12) applies and the bias factor can be evaluated using $\pi_0^* = \Phi(\mu + \beta_1^T X_1 + \xi)$, with X_1 fixed at an arbitrary value and expectations taken over the distribution of X_2 conditional on X_1. The denominator of $H'(0)$ is $\phi(h/c)$, with h as in (15) and the numerator is $\mathsf{E}[\phi(\mu + \beta_1^T X_1 + \xi)]$. This last expectation has an analytic solution leading to

$$H'(0) = \frac{\mathsf{E}(\phi[\Phi^{-1}(\pi_0)])}{\phi(\Phi^{-1}[\mathsf{E}(\pi_0)])} = \frac{\frac{\phi(h/c)}{\sqrt{1+\mathsf{var}(\xi)}}}{\phi(h/c)} = \frac{1}{\sqrt{1 + \beta_2^T \tilde{\Omega} \beta_2}}.$$

This coincides with the result from [2] for small β_2 and is the same correction factor as obtained from the use of the skew-Normal distribution. The derivation in [10] assumes that α is small and our derivation of the above expression has assumed that the covariates have a multivariate Normal distribution. In all cases the bias correction for probit regression, unlike logistic regression, depends only on the conditional variance of the omitted variables and their associated regression coefficients, and not on any measure of location.

4 Some numerical results

4.1 Assessment of the accuracy of the approximations

The simulation results in Table 1 assess the accuracy of the forms of α^* for the logistic regression given in equations (10), (11) and (14), with the last adapted as in (15) as necessary. The simulated value is found by fitting the reduced model to a sample of size 2×10^6 simulated from the full model: all calculations were performed in R , version 3.10 [18]. Three cases are presented: in the first the true model has two Normal covariates, neither of which is fitted, while in the second model only one of these covariates is omitted. The third model has five covariates, three of which are omitted. In all cases the Normal covariates have mean 0 and unit variance and correlations are 0.5. The treatment effect measured by α is taken to be 0.5. We initially take $\beta_k = 0.5$, for $k = 1, 2$ in the first cases and $1, \ldots, 5$ in the final case. It is important that the simulations correspond to realistic models, with outcome

probabilities taking values that are appropriate for a clinical trial. From (1) we find that

$$\Pr(Y = 1 \mid T = \pm 1) \approx \Phi\left(\frac{c(\mu \pm \alpha + \beta^T \nu)}{\sqrt{1 + c^2 \beta^T \Omega \beta}}\right),$$

so if μ is chosen so that $\mu + \beta^T \nu = 0$ then the outcome probabilities will be around 0.5.

Table 1 shows that when $\alpha = 0.5$ and $\beta_k = 0.5$, all methods perform reasonably when no Normal covariates are fitted, with that from (14) doing best. When some Normal covariates are fitted, Gail's method is not applicable but the proposed extension to (14) does well. The method based on the skew-Normal approximation is conservative, as would be predicted from Figure 1 for response probabilities around 0.5. When the β_k are larger Gail's method fails, as would be anticipated from its derivation. The method due to Neuhaus and colleagues performs better than the skew-Normal factor when no Normal covariates are fitted, but the skew-Normal does better when the fitted model contains some Normal covariates. The method leading to (14) assumes α is small and the final part of Table 1 shows that for large α the skew-Normal approximation is again better when Normal covariates are fitted and performs better relative to the method of Neuhaus *et al.* than it did for the smaller value of α.

When probit regression is used, the skew-Normal and Neuhaus *et al.* expressions coincide and are very close to the simulated value of α^*, across a range of values of μ: see Table 2. Gail's version is reasonable for small α and β_k but is poor for larger β_k. The lack of dependence of the corrections on μ is confirmed by the simulated α^*, which changes little with μ. This contrasts with the situation for logistic regression where the simulated α^* show that the bias reduces as $|\mu|$ increases. The phenomenon applies for all cases but is most clearly seen for small β_k and when no Normal covariates are fitted. This difference between logistic and probit regressions does not appear to be widely appreciated.

4.2 Assessment of the effect of departures from Normality

Some simulations were carried out to assess the effect of non-Normality on the performance of the expressions for α^* in (10) and (14). Two types of departure were considered. The effect of a symmetric non-Normal distribution was assessed by generating X from a central multivariate t-distribution with 4 degrees of freedom, while the effect of skewness was assessed using

	$p=0,\ q=2$	$p=1,\ q=1$	$p=2,\ q=3$
	$\alpha=0.5\ \beta_k=0.5$		
Numerical	0.433	0.482	0.308
Skew-Normal	0.446	0.485	0.330
Gail's method	0.408	-	-
Neuhaus *et al.*	0.434	0.481	0.309
	$\alpha=1.5\ \beta_k=0.5$		
Numerical	1.307	1.447	1.328
Skew-Normal	1.337	1.454	1.337
Gail's method	1.262	-	-
Neuhaus *et al.*	1.302	1.442	1.302
	$\alpha=0.5\ \beta_k=2$		
Numerical	0.206	0.347	0.227
Skew-Normal	0.220	0.350	0.220
Gail's method	-0.970	-	-
Neuhaus *et al.*	0.202	0.330	0.202
	$\alpha=1.5\ \beta_k=2$		
Numerical	0.619	1.045	0.677
Skew-Normal	0.661	1.051	0.661
Gail's method	-2.311	-	-
Neuhaus *et al.*	0.605	0.990	0.605

Table 1: Values of α^* computed using simulation (sample of size 2×10^6) and the three approximations given in equations (10), (11) and (14), for various values of the regression parameters. The Normal covariates have mean 0, unit variance and pairwise correlation of $\frac{1}{2}$. The number of fitted Normal covariates is p and the number omitted is q: throughout $\mu=0$.

the log-Normal distribution. In the latter case X was derived from a bivariate Normal variable W with zero mean. To assess the effect of skewness in the fitted or omitted variable or both, three types of model were considered, with (X_1, X_2) taken as, respectively, $(\exp(W_1)', W_2)$, $(W_1, \exp(W_2)')$ and $(\exp(W_1)', \exp(W_2)')$, where as usual X_1 is the fitted covariate and X_2 is omitted and $'$ denotes centring to zero mean. The parameters of the t and log-Normal distributions were chosen to give X_1, X_2 unit variance and correlation close to $\frac{1}{2}$, which implies that the skewness the log-Normal variables are 2.84. In all simulations $\mu=0$, with $\beta_k=0.5$ or 2 and $\alpha=0.5$ or 1.5, and one scalar covariate is fitted and one omitted. The correction factors $\tilde{q}^{-1}$ and $T(h,a)/T(h,1)$ both depend solely on the mean and variance of the X_is, so

	$p=0,\ q=2;\ \ \alpha=0.5,\ \beta_k=0.5$					
	Logistic regression			Probit regression		
	$\mu=0$	$\mu=2$	$\mu=4$	$\mu=0$	$\mu=2$	$\mu=4$
Numerical	0.433	0.455	0.488	0.378	0.377	0.372
Skew-Normal	0.446	0.446	0.446	0.378	0.378	0.378
Gail's method	0.408	0.460	0.493	0.313	0.313	0.313
Neuhaus *et al.*	0.434	0.452	0.482	0.378	0.378	0.378
	$p=0,\ q=2;\ \ \alpha=0.5,\ \beta_k=2$					
	Logistic regression			Probit regression		
	$\mu=0$	$\mu=2$	$\mu=4$	$\mu=0$	$\mu=2$	$\mu=4$
Numerical	0.207	0.213	0.231	0.139	0.138	0.140
Skew-Normal	0.220	0.220	0.220	0.139	0.139	0.139
Gail's method	-0.971	-0.139	0.390	-2.50	-2.50	-2.50
Neuhaus *et al.*	0.202	0.208	0.227	0.139	0.139	0.139
	$p=2,\ q=3;\ \ \alpha=0.5,\ \beta_k=0.5$					
	Logistic regression			Probit regression		
	$\mu=0$	$\mu=2$	$\mu=4$	$\mu=0$	$\mu=2$	$\mu=4$
Numerical	0.437	0.445	0.458	0.378	0.378	0.379
Skew-Normal	0.446	0.446	0.446	0.378	0.378	0.378
Neuhaus *et al.*	0.434	0.452	0.482	0.378	0.378	0.378

Table 2: Values of α^* computed using simulation (sample of size 2×10^6) and the approximations, for both logistic and probit regression, for different locations of the linear predictor. The Normal covariates have mean 0, unit variance and pairwise correlation of $\frac{1}{2}$. The number of fitted Normal covariates is p and the number omitted is q.

these will be the same for all of the above models.

From Table 3 we see that for smaller β_k the predictions of bias provided by (10) and (14) remain accurate even when the covariates have non-Normal distributions. For larger values of β_k, α^* tends to be closer to α for these non-Normal covariates than for Normal covariates. However, it should be noted that in this context $\beta_k = 2$ is a large coefficient for a covariate with unit variance and unlikely to be encountered in practice.

	$\alpha = 0.5$ $\beta_k = 0.5$	$\alpha = 1.5$ $\beta_k = 0.5$	$\alpha = 0.5$ $\beta_k = 2$	$\alpha = 1.5$ $\beta_k = 2$
X bivariate t, 4 df	0.484	1.456	0.376	1.129
$X = (\exp(W_1), W_2)$	0.479	1.441	0.352	1.061
$X = (W_1, \exp(W_2))$	0.488	1.460	0.403	1.194
$X = (\exp(W_1), \exp(W_2))$	0.481	1.452	0.375	1.131
Skew-Normal	0.485	1.454	0.350	1.051
Neuhaus *et al.*	0.481	1.442	0.330	0.990

Table 3: Values of α^* computed using simulation (sample of size 2×10^6) and the two approximations given in equations (10) and (14), for various values of the regression parameters. The covariates have a multivariate t distribution with 4 df or are a mixture of Normal and log-Normal variables. In each case one covariate is fitted and one omitted, in addition to the treatment indicator: throughout $\mu = 0$. The approximations below the line apply to all the cases above it.

4.3 An example: the Mayo Clinic primary biliary cirrhosis trial

No direct evaluation of the above results is possible as they are all expressed in terms of parameter values. However, some practical indication of the size of the asymptotic bias, and how this changes with the included covariates, would be helpful. A trial with binary outcome and several Normal baseline covariates is the primary biliary cirrhosis trial conducted at the Mayo Clinic over ten years from 1974: the trial randomized patients to placebo or penicillamine [19], and the data are given in [20]. By way of illustration we take end-of-study mortality as the outcome and fit a model with a treatment indicator and five continuous baseline covariates, namely the serum values of bilirubin (mg/dl), cholesterol (mg/dl), albumin (gm/dl), urinary copper (μg/day) and alkaline phosphatase(AP) (U/litre). All variables but albumin were log-transformed (base 10) to achieve Normality.

The dispersion matrix of the five baseline covariates, based on the 312 patients in the trial, was used as Ω and β was taken to be the estimated regression coefficients from the full logistic regression. The values of $\sqrt{1 + c^2 \beta_a^T \tilde{\Omega} \beta_a}$ were then computed for a sequence of models in which the first model includes only the treatment indicator, the second also includes log bilirubin, and then, successively, log cholesterol, albumin and log copper are added. The correlations are shown in Table 4 and the $\tilde{q}$ values are in Table 5.

If we assume that the model with treatment indicator and all five vari-

	β_k	log bilirubin	log cholesterol	albumin	log copper	log AP
log bilirubin	2.19	*0.201*				
log cholesterol	-1.43	0.488	*0.036*			
albumin	-0.55	-0.360	-0.038	*0.176*		
log copper	0.88	0.598	0.217	-0.278	*0.128*	
log AP	1.69	0.295	0.351	-0.146	0.277	*0.098*

Table 4: The correlations obtained from the dispersion matrix for the five continuous covariates chosen from the PBC trial, with variances on the diagonal and regression coefficients in the second column

Included variables	$\tilde{q}$
None	1.311
+ log bilirubin	1.072
+ log cholesterol	1.068
+ albumin	1.056
+ log copper	1.039

Table 5: The values of $\tilde{q} = \sqrt{1 + c^2 \beta_a^T \tilde{\Omega} \beta_a}$ for a series of increasing models

ables is the correct model, then $\hat{\alpha}$ from this model will be asymptotically unbiassed. However, if a model with no covariates is fitted, $\hat{\alpha}$ will tend to $\alpha/\tilde{q} \approx \alpha/1.3$, i.e. a value about 75% of the correct value. Including log bilirubin reduces the bias and $\hat{\alpha}$ will tend to $\alpha/1.07$, a value in error by approximately 7%. As Table 5 shows, this can be reduced further by including more covariates, although the change is never as marked as when the first variable was introduced. Of course, different results would be obtained if terms were added in a different order.

5 Extensions of the model

The analysis presented thus far applies to a model where, apart from a binary treatment indicator, the covariates are assumed to be continuous. It is often the case that in clinical trials some baseline variables are categorical. While such variables may have more than two categories, they would usually be included in a linear predictor through dummy variables, so there is no loss in assuming that categorical covariates are binary. The values of the binary treatment indicator are assigned by randomization, so are independent of

the values of the other covariates, a feature that would not be shared by a general binary covariate.

If the model in (1) were extended to include a single non-treatment binary covariate, $B \in \{-1, 1\}$, as in

$$\Pr(Y = 1 \mid T, B, X_f, X_a) = \text{expit}(\mu + \alpha T + \gamma B + \beta_1^T X_1 + \beta_2^T X_2), \quad (17)$$

then the foregoing analysis of the effect of omitting X_2 from the fitted model can be adapted to this case. In this model, as in Section 2, T is a binary indicator of the randomized treatment, so is independent of B and X. Consequently the parameters defining the distributions of B and X are unaffected by the value of T and we take $\Pr(B = b) = \theta_b$ and $\mathsf{E}(X \mid B = b) = \nu_b$, $b = -1, 1$, but continue to assume that the variance is unaffected by the value of B, i.e. $\mathsf{var}(X \mid B = b) = \Omega$.

Under these assumptions it follows that β_1^* is as in (7) and

$$\begin{aligned}
\alpha^* &\approx \frac{\alpha}{\sqrt{1 + c^2 \beta_2^T \tilde{\Omega} \beta_2}} \\
\gamma^* &\approx \frac{\gamma + \frac{1}{2}\beta_2^T([\nu_{1,a} - \nu_{-1,a}] - \Omega_{21}\Omega_{11}^{-1}[\nu_{1,1} - \nu_{-1,1}])}{\sqrt{1 + c^2 \beta_2 2^T \tilde{\Omega} \beta_2}} \\
\mu^* &\approx \frac{\mu + \frac{1}{2}\beta_2^T((\nu_{1,2} + \nu_{-1,2}) - \Omega_{21}\Omega_{11}^{-1}(\nu_{1,1} + \nu_{-1,1}))}{\sqrt{1 + c^2 \beta_2^T \tilde{\Omega} \beta_2}}.
\end{aligned}$$

where $\nu_{b,1}, \nu_{b,2}$ is the partition of ν_b corresponding to the partition of X into X_1 andX_2. The above results are exact for probit regression, provided that the factor c^2 is omitted from the denominator.

The above argument can be extended to an arbitrary number of binary covariates, $B_1, \ldots, B_K$ but only at the expense of rather restrictive assumptions about the form of $\mathsf{E}(X \mid B_1, \ldots, B_K)$.

6 Discussion

One of the main reasons for advocating the use of baseline variables in the analysis of a randomized controlled trial is to correct for treatment imbalances in these variables that arise notwithstanding the random allocation. However, the bias discussed in the present paper is because of the geometrical structure of the statistical model [10], and would apply even if the groups were perfectly balanced. If there are important covariates in the model then

omitting all of them can lead to noticeable biasses, as seen from Table 5, where the log OR is reduced by about 25%.

Failing to include important covariates can therefore have important consequences for the analysis of a trial. Not only might the effect of the treatment be underestimated, power calculations may be compromised because observed odds ratios represent less than the true treatment effect. Of course, in practice the important covariates are not known, at least not with certainty. If a covariate is included because its importance is suspected, but in fact this view is mistaken, then the corresponding element of β_1 is zero and there is no penalty in terms of the asymptotic bias in α. However, in practical applications, issues related to the finite sample size need to be taken into account. Adding extra variables to a logistic regression may reduce the bias in the treatment estimator but at the expense of an increase in its variance: a fuller investigation of this aspect of the problem is important but beyond the scope of this paper. Nevertheless, if sufficient data are available when the trial is being planned then it may be possible to use results such as (10) to help the triallist make an informed judgment about which covariates ought to be included in the final analysis.

The comparison between logistic and probit analyses is interesting. If the parameter estimates from a logistic regression are $\hat{\beta}$ then the estimates obtained from fitting a probit regression to the same data will be approximately $c\hat{\beta}$, so the corrections in (10) and (16) are essentially equal. However, as Figure 1 shows, the correction in (14), which can be more accurate than (10) for logistic regression, indicates that the asymptotic bias α^* can be greater than is implied by (10). However, the correction in (16) is an exact result, so it may be that the problem of asymptotic bias in the estimates of the treatment effect are less if probit is preferred to logistic regression.

Appendix

Least false values for logistic regression

Applying the approximation $\text{expit}(u) \approx \Phi(cu)$ to (3), (4) and (5) and using the properties of the ESN distribution, we obtain from (3) and (4) the equations

$$p_1\Phi(\psi_1^*) \pm p_{-1}\Phi(\psi_{-1}^*) = p_1\Phi(\psi_1) \pm p_{-1}\Phi(\psi_{-1}) \tag{18}$$

and from (5) we obtain

$$p_1[\nu_1\Phi(\psi_1^*) + \frac{c\Omega_{11}\beta_1^*}{\sqrt{1+c^2\beta_1^{*T}\Omega_{11}\beta_1^*}}\phi(\psi_1^*)] + p_{-1}[\nu_{-1}\Phi(\psi_{-1}^*) + \frac{c\Omega_{11}\beta_1^*}{\sqrt{1+c^2\beta_1^{*T}\Omega_{11}\beta_1^*}}\phi(\psi_{-1}^*)]$$
$$= p_1[\nu_1\Phi(\psi_1) + \frac{c(\Omega\beta)_1}{\sqrt{1+c^2\beta^T\Omega\beta}}\phi(\psi_1)] + p_{-1}[\nu_{-1}\Phi(\psi_{-1}) + \frac{c(\Omega\beta)_1}{\sqrt{1+c^2\beta^T\Omega\beta}}\phi(\psi_{-1})] \tag{19}$$

Here $p_t = \Pr(T=t)$, $\beta^T = (\beta_1^T, \beta_2^T)^T$, $(\Omega\beta)_1$ denotes the first p elements of $\Omega\beta$ and

$$\psi_1^* = \frac{c(\mu_1^* + \alpha^*)}{\sqrt{1+c^2\beta_1^{*T}\Omega_{11}\beta_1^*}} \qquad \psi_{-1}^* = \frac{c(\mu_{-1}^* - \alpha^*)}{\sqrt{1+c^2\beta_1^{*T}\Omega_{11}\beta_1^*}}$$
$$\psi_1 = \frac{c(\mu_1 + \alpha)}{\sqrt{1+c^2\beta^T\Omega\beta}} \qquad \psi_{-1} = \frac{c(\mu_{-1} - \alpha)}{\sqrt{1+c^2\beta^T\Omega\beta}}$$

with $\mu_t^* = \mu^* + \beta_1^{*T}\nu_{t,1}$ and $\mu_t = \mu + \beta^T\nu_t$, where $\nu_{t,1}$ is written for the first p elements of ν_t. From (18) we obtain $\psi_1^* = \psi_1$ and $\psi_{-1}^* = \psi_{-1}$, and using this in (19) we get

$$\frac{\Omega_{11}\beta_1^*}{\sqrt{1+c^2\beta_1^{*T}\Omega_{11}\beta_1^*}} = \frac{(\Omega\beta)_1}{\sqrt{1+c^2\beta^T\Omega\beta}}$$

and these can be solved to give (7), (8) and (9).

Least false values for probit regression

The least false equations for the maximum likelihood estimators for a probit regression differ from (3) to (5) because of the presence of a weighting factor $\omega = \omega(T, X_1) = \omega(\eta^*)$ with $\eta^* = \mu^* + \alpha^* T + \beta_1^{*T} X_1$, i.e. the $p+2$ equations

$$\mathsf{E}[\omega(\eta^*)Z\Phi(\eta^*)] = \mathsf{E}[\omega(\eta^*)Z\Phi(\mu + \alpha T + \beta_1^T X_1 + \beta_2^T X_2)] \tag{20}$$

where $\omega(\eta^*) = \phi(\eta^*)/[\Phi(\eta^*)\Phi(-\eta^*)]$, and where Z is taken to be, successively, 1, T and X_{1j}, $j = 1, \ldots, p$. The presence of ω means that the skew-Normal distribution cannot be used to evaluate the expectations in the way it was used for logistic regression, but it can be applied to evaluate the right hand expectation in (20) over the distribution of X_2 conditional on T and X_1, giving

$$\mathsf{E}\left[\omega(\eta^*)Z\Phi\left\{\frac{\mu + \beta_2^T(\nu_{T,2} - \Omega_{21}\Omega_{11}^{-1}\nu_{T,1}) + \alpha T + (\beta_1 + \Omega_{11}^{-1}\Omega_{12}\beta_2)^T X_1}{\sqrt{1+\beta_2^T\tilde{\Omega}\beta_2}}\right\}\right]$$

Consequently, if we choose β_1^*, μ^* and α^* as in (7), (8) and (9) but with denominator $\sqrt{1+\beta_2^T\tilde{\Omega}\beta_2}$ as opposed to $\sqrt{1+c^2\beta_2^T\tilde{\Omega}\beta_2}$, then equations (20) will be satisfied.

References

[1] Cox DR. *Planning of Experiments*. Wiley Classics Library edn., Wiley: Chichester, 1992.

[2] Gail MH, Wieand S, Piantadosi S. Biased estimates of treatment effect in randomized experiments with nonlinear regressions and omitted covariates. *Biometrika* 1984; **71**:431–444.

[3] Lee LF. Specification error in multinomial logit models. *Journal of Econometrics* 1982; **20**:197–209.

[4] Hauck WW, Neuhaus JM, Kalbfleisch JD, Anderson S. A consequence of omitted covariates when estimating odds ratios. *Journal of Clinical Epidemiology* 1991; **44**:77–81.

[5] Robinson LD, Jewell NP. Some surprising results about covariate adjustment in logistic regression models. *International Statistical Review* 1991; **59**:227–240.

[6] melissa Dowd Begg, Lagakos S. Loss in efficiency caused by omitting covariates and misspecifying exposure in logistic regression models. *Journal of the American Statistical Association* 1993; **88**:166–170.

[7] Lin DY, Psaty BM, Kronmal RA. Assessing the sensitivity of regression results to unmeasured confounders in observational studies. *Biometrics* 1998; **54**:948–963.

[8] Drake C, McQuarrie A. A note on the bias due to omitted confounders. *Biometrika* 1995; **82**:633–638.

[9] Neuhaus JM, Kalbfleisch JD, Hauck WW. A comparison of cluster-specific and population-averaged approaches for analyzing correlated binary data. *International Statistical Review* 1991; **59**:25–35.

[10] Neuhaus JM, Jewell NP. A geometric approach to assess bias due to omitted covariates in generalized linear models. *Biometrika* 1993; **80**:807–815.

[11] Hauck WW, Anderson S, Marcus SM. Should we adjust for covaraites in nonlinear regression analyses of randomized trials? *Controlled Clinical Trials* 1998; **19**:249–256.

[12] Arnold BC, Beaver RJ. Hidden truncation models. *Sankhyā A* 2000; **62**:23–35.

[13] Claeskens G, Hjort NL. *Model Selection and Model Averaging*. IMS Monographs, Cambridge University Press: Cambridge, 2008.

[14] Johnson NL, Kotz S, Balakrishnan N. *Continuous Univariate Distributions, volume 2*. 2nd edn., Wiley: Chichester, 1995.

[15] Azzalini A, Capitanio A. *The Skew-Normal and Related Families*. Cambridge University Press: Cambridge, 2014.

[16] Owen DB. Tables for computing bivariate normal probabilities. *Annals of Mathematical Statistics* 1956; **27**:1075–1090.

[17] Azzalini A. *The R* `sn` *package : The skew-normal and skew-t distributions (version 1.0-0)*. Università di Padova, Italia 2014. URL `http://azzalini.stat.unipd.it/SN`.

[18] R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria 2014. URL `http://www.R-project.org`.

[19] Dickson ER, Fleming TR, Wiesner RH, Baldus WP, Fleming CR, Ludwig J, McCall JT. Trial of penicillamine in advanced primary billiary cirrhosis. *New England Journal of Medicine* 1985; **312**:1011–1015.

[20] Fleming TR, Harrington DP. *Counting Processes and Survival Analysis*. 2nd edn., Wiley: Chichester, 2005.